



Introduction to image processing

The ESA/ESO/NASA Photoshop FITS Liberator v.2

1. Introduction to image processing

1.1 What is an image?

An image is an array, or a matrix, of square pixels (picture elements) arranged in columns and rows.

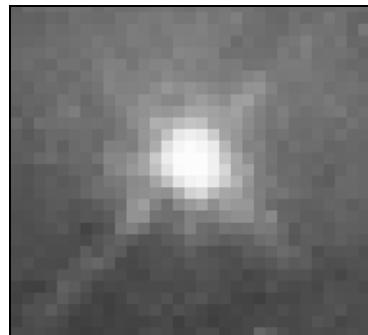


Figure 1: An image – an array or a matrix of pixels arranged in columns and rows.

In a (8-bit) greyscale image each picture element has an assigned intensity that ranges from 0 to 255. A grey scale image is what people normally call a black and white image, but the name emphasizes that such an image will also include many shades of grey.

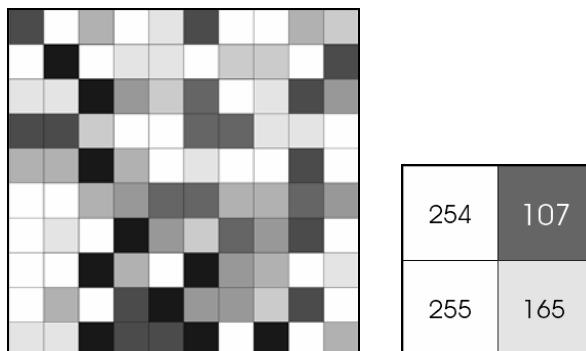


Figure 2: Each pixel has a value from 0 (black) to 255 (white). The possible range of the pixel values depend on the colour depth of the image, here 8 bit = 256 tones or greyscales.

A normal greyscale image has 8 bit colour depth = 256 greyscales. A “true colour” image has 24 bit colour depth = $8 \times 8 \times 8$ bits = $256 \times 256 \times 256$ colours = ~ 16 million colours.

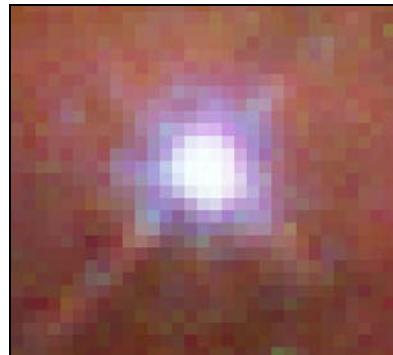


Figure 3: A true-colour image assembled from three greyscale images coloured red, green and blue. Such an image may contain up to 16 million different colours.

Some greyscale images have more greyscales, for instance 16 bit = 65536 greyscales. In principle three greyscale images can be combined to form an image with 281,474,976,710,656 greyscales.

There are two general groups of ‘images’: vector graphics (or line art) and bitmaps (pixel-based or ‘images’). Some of the most common file formats are:

- GIF — an 8-bit (256 colour), non-destructively compressed bitmap format. Mostly used for web. Has several sub-standards one of which is the animated GIF.
- JPEG — a very efficient (i.e. much information per byte) destructively compressed 24 bit (16 million colours) bitmap format. Widely used, especially for web and Internet (bandwidth-limited).
- TIFF — the standard 24 bit publication bitmap format. Compresses non-destructively with, for instance, Lempel-Ziv-Welch (LZW) compression.
- PS — Postscript, a standard vector format. Has numerous sub-standards and can be difficult to transport across platforms and operating systems.
- PSD – a dedicated Photoshop format that keeps all the information in an image including all the layers.

1.2 Colours

For science communication, the two main colour spaces are RGB and CMYK.

1.2.1 RGB

The RGB colour model relates very closely to the way we perceive colour with the **r**, **g** and **b** receptors in our retinas. RGB uses additive colour mixing and is the basic colour model used in television or any other medium that projects colour with light. It is the basic colour model used in computers and for web graphics, but it cannot be used for print production.

The secondary colours of RGB – cyan, magenta, and yellow – are formed by mixing two of the primary colours (red, green or blue) and excluding the third colour. Red and green combine to make yellow, green and blue to make cyan, and blue and red form magenta. The combination of red, green, and blue in full intensity makes white.

In Photoshop using the “screen” mode for the different layers in an image will make the intensities mix together according to the additive colour mixing model. This is analogous to stacking slide images on top of each other and shining light through them.

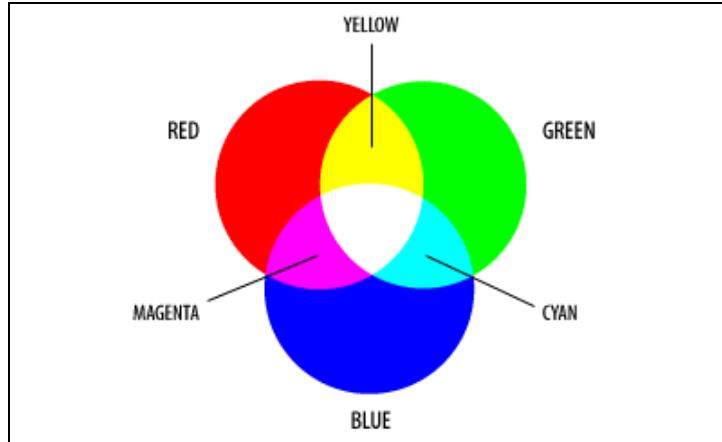


Figure 4: The additive model of RGB. Red, green, and blue are the primary stimuli for human colour perception and are the primary additive colours. Courtesy of adobe.com.

1.2.2 CMYK

The 4-colour CMYK model used in printing lays down overlapping layers of varying percentages of transparent cyan (C), magenta (M) and yellow (Y) inks. In addition a layer of black (K) ink can be added. The CMYK model uses the subtractive colour model.

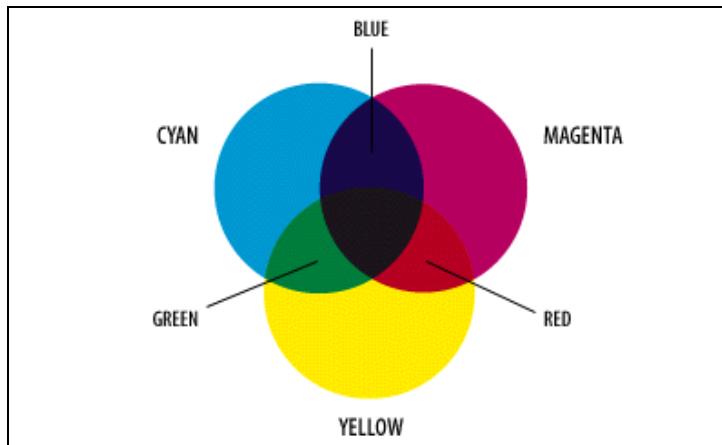


Figure 5: The colours created by the subtractive model of CMYK don't look exactly like the colours created in the additive model of RGB. Most importantly, CMYK cannot reproduce the brightness of RGB colours. In addition, the CMYK gamut is much smaller than the RGB gamut. Courtesy of adobe.com.

1.2.3 Gamut

The range, or gamut, of human colour perception is quite large. The two colour spaces discussed here span only a fraction of the colours we can see. Furthermore the two spaces do not have the same gamut, meaning that converting from one colour space to the other may cause problems for colours in the outer regions of the gamuts.

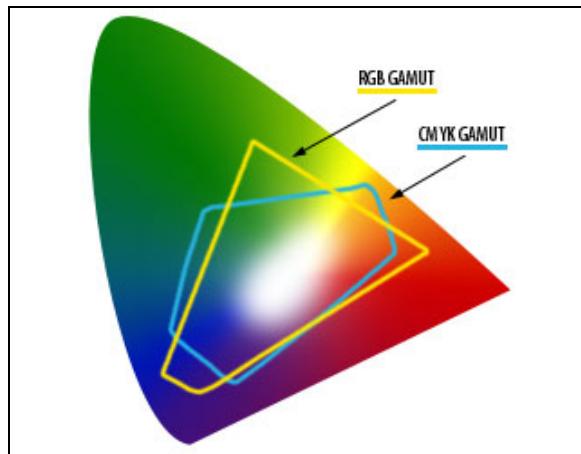


Figure 6: This illustration clearly shows the different gamuts of the RGB and CMYK colour spaces. The background is the CIE Chromaticity Diagram (representing the whole gamut of human colour perception). Courtesy adobe.com.

1.3 Astronomical images

Images of astronomical objects are usually taken with electronic detectors such as a CCD (Charge Coupled Device). Similar detectors are found in normal digital cameras. Telescope images are nearly always greyscale, but nevertheless contain some colour information. An astronomical image may be taken through a colour filter. Different detectors and telescopes also usually have different sensitivities to different colours (wavelengths).

1.3.1 Filters

A telescope such as the NASA/ESA Hubble Space Telescope typically has a fixed number of well-defined filters. A filter list for Hubble's WFPC2 (Wide Field and Planetary Camera 2) camera is seen below.



Name	Type	Wheel	Slo t	Notes	In WF/PC-1?	$\bar{\lambda}$ (Å)	$\Delta\bar{\lambda}$ (Å)	Peak T (%)	Peak λ (Å)
F122M	A	1	4	H Ly α - Red Leak	Y	1259	224.4	19.3	1240
F130LP	B	2	1	CaF2 Blocker (zero focus)	N	2681	5568.3	94.5	8852
F160AW	A	1	3	Woods A - redleak from pinholes	N	1471	457.2	10.1	1403
F160BW	A	1	2	Woods B	N	1446	457.1	12.1	1400
F165LP	B	2	2	Suprasil Blocker (zero focus)	N	3301	5533.2	95.4	5796
F170W	A	8	1	-	N	1666	434.6	30.7	1655
F185W	A	8	2	-	N	1899	297.4	25.9	1849
F218W	A	8	3	Interstellar feature	N	2117	367.9	21.1	2092
F255W	A	8	4	-	N	2545	408.2	14.8	2489
F300W	A	9	4	Wide U	N	2892	727.6	50.8	2760
F336W	A	3	1	WFPC2 U, Strömgren u	Y	3317	370.5	82.6	3447
F343N	A	5	1	Ne V	N	3427	23.5	9.3	3432
F375N	A	5	2	[OII] 3727 RS	Y	3732	24.4	19.5	3736
F380W	A	9	1	-	N	3912	694.8	65.0	3980
F390N	A	5	3	CN	N	3888	45.0	36.5	3886
F410M	A	3	2	Strömgren v	N	4086	147.0	70.4	4097
F437N	A	5	4	[OIII]	Y	4369	25.2	52.0	4368
F439W	A	4	4	WFPC2 B	Y	4283	464.4	68.2	4176
F450W	A	10	4	Wide B	N	4410	925.1	91.4	5060
F467M	A	3	3	Strömgren b	N	4663	166.4	75.3	4728
F469N	A	6	1	He II	Y	4694	25.0	52.4	4697
F487N	A	6	2	H β	Y	4865	25.9	58.6	4862
F502N	A	6	3	[OIII]	Y	5012	26.9	63.7	5008
F547M	A	3	4	Strömgren y (but wider)	Y	5446	486.6	91.3	5360
F555W	A	9	2	WFPC2 V	Y	5202	1222.6	94.6	5148
F569W	A	4	2	F555W generally preferred ^a	Y	5524	965.7	94.2	5310
F588N	A	6	4	He I & Na I (NaD)	Y	5893	49.0	91.4	5894
F606W	A	10	2	Wide V	Y	5767	1579.0	96.7	6186
F622W	A	9	3	-	Y	6131	935.4	95.6	6034
F631N	A	7	1	[OI]	Y	6306	30.9	85.7	6301
F656N	A	7	2	H α	Y	6564	21.5	77.8	6562
F658N	A	7	3	[NII]	Y	6591	28.5	79.7	6591
F673N	A	7	4	[SII]	Y	6732	47.2	87.0	6732
F675W	A	4	3	WFPC2 R	Y	6714	889.5	97.3	6780
F702W	A	10	3	Wide R	Y	6940	1480.6	97.1	6538
F785LP	A	2	3	F814W generally preferred ^a	Y	9283	2096.1	91.7	9959
F791W	A	4	1	F814W generally preferred ^a	Y	7969	1304.6	95.9	8082
F814W	A	10	1	WFPC2 I	Y	8203	1758.0	94.8	8387
F850LP	A	2	4	-	Y	9650	1672.4	89.2	10028
F953N	A	1	1	[SIII]	N	9546	52.5	95.6	9528
F1042M	A	11	2	-	Y	10437	611.0	81.6	10139

Figure 7: Filter list for Hubble's WFPC2 camera (Wide Field and Planetary Camera 2). Filter names are to the left (names include approximate wavelength in nm) in column 1. Column 5 contains the physical property of the radiation the filter lets through. Column 7 is the central wavelength. The N's and W's are short for Narrow and Wide.

Filters can either be broad-band (**Wide**) or narrow-band (**Narrow**). A broad-band filter lets a wide range of colours through, for instance the entire green or red area of the spectrum. A narrow-band filter typically only lets a small wavelength span through, thus effectively restricting the transmitted radiation to that coming from a given atomic transition, allowing astronomers to investigate individual atomic processes in the object.

A filename such as 502nmos.fits indicates that the filter used has a peak at 502 nm. In the table below, you can see that this filter is a narrow bandwidth filter, i.e. it only lets radiation with wavelengths within a few nm of 502 nm through.



Below is an example of an image composed from narrow-band exposures. This results in very sharply defined wisps of nebulosity since each exposure separates light from only some very specific physical processes and locations in the nebula.



Figure 7: Example of an image constructed from narrow-band exposures. Since the narrow-band exposures probe individual atomic transitions the result is an image that has very 'sharp' features.

Galaxies are often studied through broad-band filters as they allow more light to get through. Also the processes in a galaxy are more 'mixed' or complicated, result from the outputs of billions of stars and so narrow-band filters give less 'specific' information about the processes there.

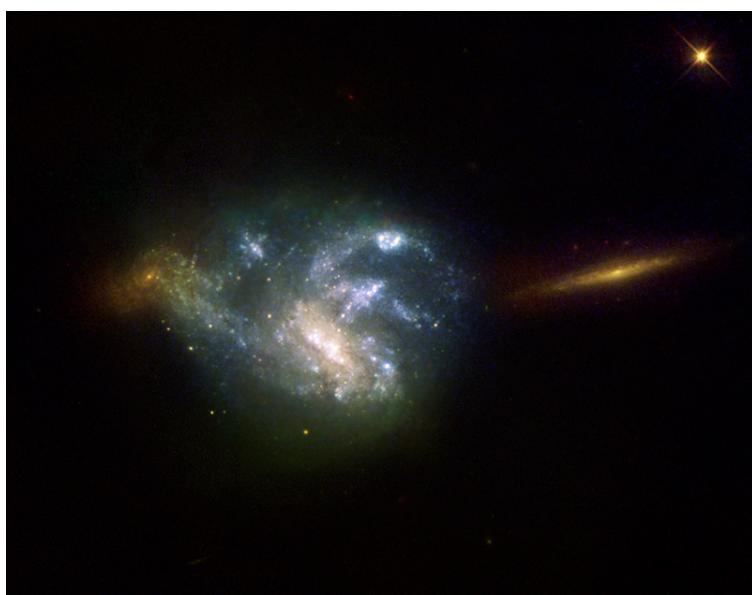


Figure 8: A broad-band image of the "Hyperactive galaxy NGC 7673".



A visual example of the different filters available onboard Hubble is seen in the following figure.

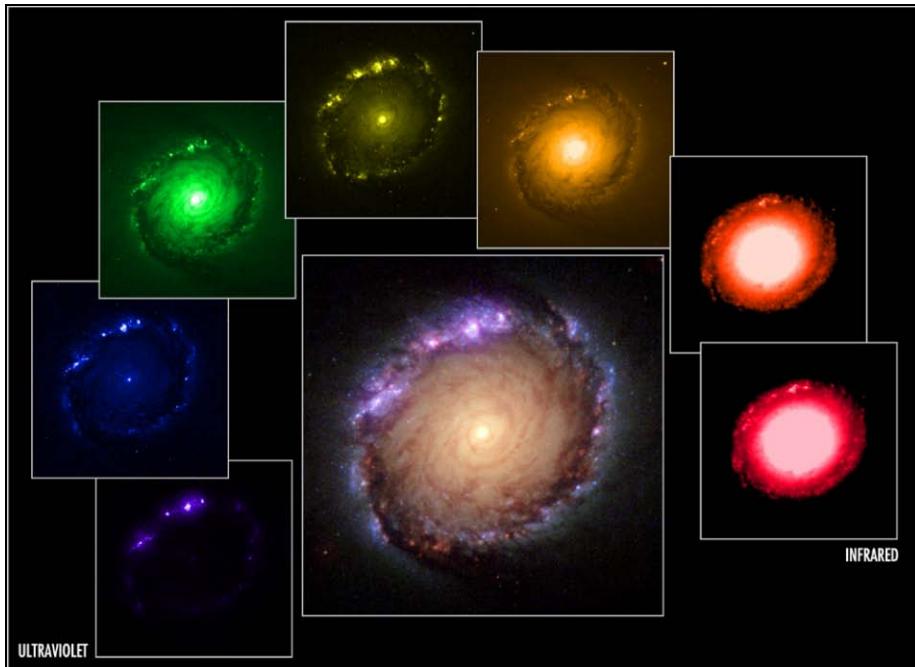


Figure 9: An example of an image constructed from 7 broad-band filters all the way from ultraviolet (left) to infrared (right).

A figure illustrating the process of stacking together different colour exposures is seen below.

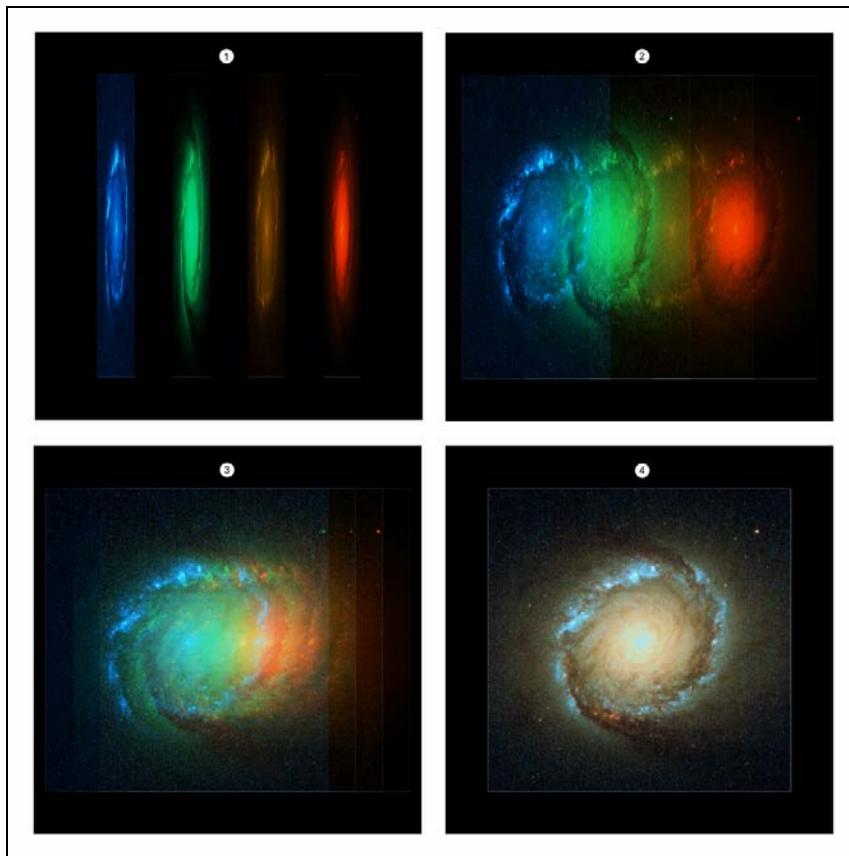


Figure 10: An example of how a colour image is constructed from four broad-band filters (seen from the side in 1.): blue, green yellow and red. When the images are overlaid (2. and 3.) the resulting image (4.) is a colour composite.

1.4 Assigning colours to different filter exposures

The astronomical images we see on the web and in the media are usually ‘refined’ or ‘processed’ as compared to the raw data that the astronomers work on with their computers. In ‘pretty pictures’ all artefacts coming from the telescope or the detectors are for instance removed as they do not say anything about the objects themselves. It is very rare that images are taken with the sole intention of producing a ‘pretty’ colour picture. Most ‘pretty pictures’ are constructed from data that was acquired to study some physical process, and the astronomer herself probably never bothered to assemble the greyscale images to a colour image.

Natural colour images

It is possible to create colour images that are close to “true-colour” if three wide band exposures exist, and if the filters are close to the r, g and b receptors in our eyes. Images that approximate what a fictitious space traveller would see if he or she actually travelled to the object are called “natural colour” images.

To make a natural colour image the order of the colours assigned to the different exposures should be in “chromatic order”, i.e. the lowest wavelength should be given a blue hue, the middle wavelength a green hue and the highest wavelength should be red.

Representative colour images

If one or more of the images in a data set is taken through a filter that allows radiation that lies outside the human vision span to pass – i.e. it records radiation invisible to us – it is of course not possible to make a natural colour image. But it is still possible to make a colour image that shows important information about the object. This type of image is called a representative colour image. Normally one would assign colours to these exposures in chromatic order with blue assigned to the shortest wavelength, and red to the longest. In this way it is possible to make colour images from electromagnetic radiation far from the human vision area, for example x-rays. Most often it is either infrared or ultraviolet radiation that is used.

Enhanced colour images

Sometimes there are reasons to not use a chromatic order for an image. Often these reasons are purely aesthetic, as is seen in the example below. This type of colour image is called an enhanced colour image.



Figure 12: An example of an enhanced colour image (not in chromatic order): Sometimes it is necessary to break the ‘rules’ for image processing. Here the Hydrogen-alpha filter is coloured blue instead of the red colour it is in nature. This is an example of a so-called false-colour image, where the blue was chosen for aesthetic reasons.

You are the judge

When processing raw science images one of the biggest problems is that, to a large degree, you are ‘creating’ the image and this means a colossal freedom within a huge parameter space. There are literally thousands of sliders, numbers, dials, curves etc. to twist and turn.

Speaking of right and wrong, there really are no wrong or right images. There are some fundamental scientific principles that should normally be observed, but the rest is a matter of aesthetics — taste. Chromatic ordering of the exposures is one of the important scientific principles.

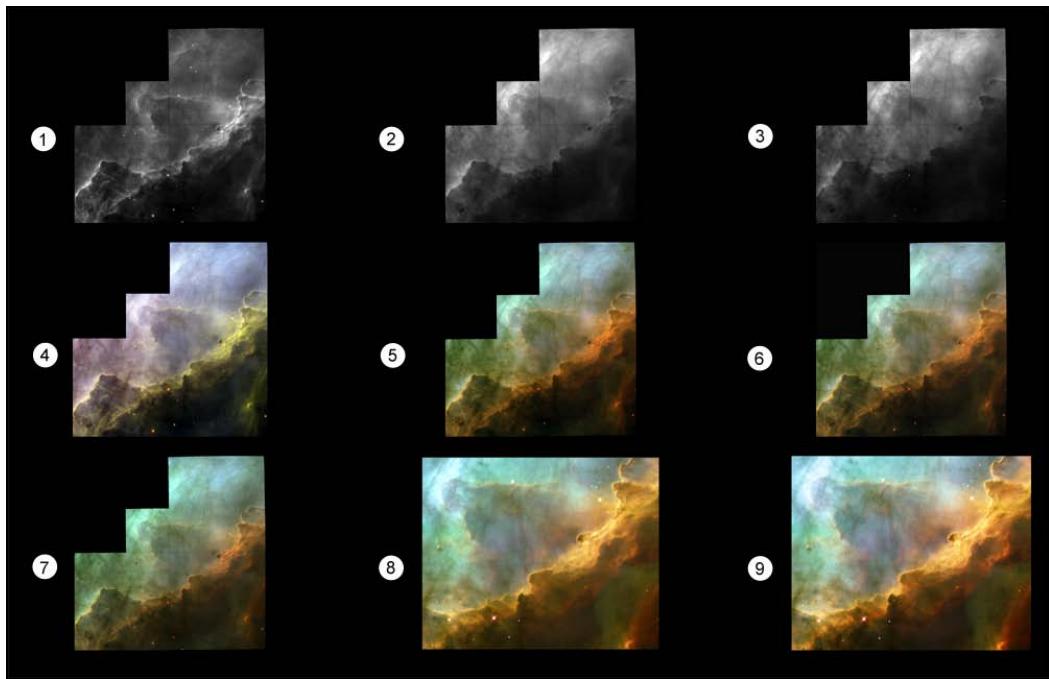


Figure 13: Sequences in the production of a Hubble Space Telescope image of Messier 17.
 First the individual exposure (taken through three different filters): 1. 673n (Sulphur) shown in red in the final image), 2. 656n (hydrogen, green), 3. 502n (oxygen, blue), 4. First colour composite attempt, 5. Improving, 6. Improving, 7. Improving, 8. Adjusting the composition and then 9. Final colour and contrast adjustments for the final image.

1.5 Stretch function

One particularly important aspect of image processing is the choice of the best stretch function. You choose which “stretch function” or representation to use in the Fits Liberator window.

A logarithmic representation of the pixel values tends to suppress the bright parts of the image, i.e. the stars, and to enhance the fainter part, e.g. nebulosity. This can be desirable if the ‘faint stuff’ needs ‘a boost’, but a logarithmic stretch function can also reduce the contrast in an image, producing a lower dynamic range as is seen in the example below.

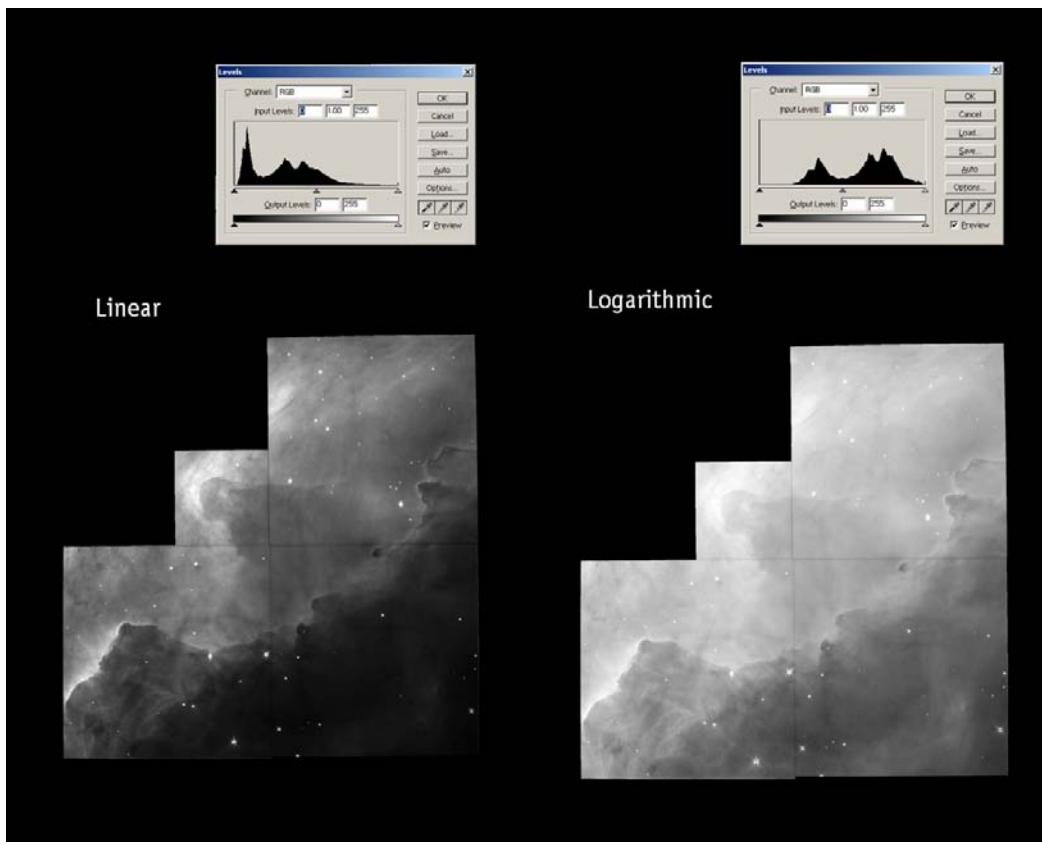


Figure 14: The difference between two stretch functions. To the left is a linear representation of the pixels and to the right a logarithmic. It is seen that the log lowers the contrast too much and therefore is not the aesthetically desirable function to choose here.

Links

http://www.hubblesite.org/sci.d.tech/behind_the_pictures/

http://heritage.stsci.edu/commonpages/infoindex/ourimages/color_comp.html